

Solving Orienteering with Category Constraints Using Prioritized Search

Paolo Bolzoni

Free University of Bozen-Bolzano
Faculty of Computer Science
Piazza Domenicani 3
39100 Bolzano; Italy

paolo.bolzoni@stud-inf.unibz.it

Sven Helmer

Free University of Bozen-Bolzano
Faculty of Computer Science
Piazza Domenicani 3
39100 Bolzano; Italy

shelmer@inf.unibz.it

ABSTRACT

We develop an approach for solving rooted orienteering problems with category constraints as found in tourist trip planning and logistics. It is based on expanding partial solutions in a systematic way, prioritizing promising ones, which reduces the search space we have to traverse during the search. The category constraints help in reducing the space we have to explore even further. We implement an algorithm that computes the optimal solution and also illustrate how our approach can be turned into an approximation algorithm, yielding much faster run times and guaranteeing lower bounds on the quality of the solution found. We demonstrate the effectiveness of our algorithms by comparing them to the state-of-the-art approach and an optimal algorithm based on dynamic programming, showing that our technique clearly outperforms these methods.

1. INTRODUCTION

Imagine a user arriving at the train station of a city, depicted on the left-hand side of Figure 1, and needing to be at the airport, located on the right-hand side of the figure, five hours later. They do not want to immediately go to the airport, but have a look at the city first. If the user takes the shortest route from the train station to the airport, represented by the solid line, they are only able to see a basilica on the way and still have a lot of spare time when arriving at the airport. According to the ratings of a tourist guide, the basilica, a pagoda, and a cathedral are the top points of interest (POIs) that can be visited en route to the airport within five hours (see dashed line). While this makes much better use of the time, the user may not be in the mood to visit these POIs, as they may be tired from traveling and would like a more relaxing route. In this instance, the dotted line, connecting a park, a ferris wheel, and a statue, is much more appropriate. As this example shows, the goal is to find a trip with the best points of interest that makes good use of the available time while at the same time considering the preferences of a user.

Finding itineraries for domains such as tourist trip planning and logistics often involves solving an orienteering problem. This is because those tasks are not about determining the shortest path, but the most attractive or the one covering the most needy customers while satisfying a strict time constraint. We focus on a variant that assumes that every point of interest (POI) or customer has a category. This categorization helps a user in expressing preferences, e.g. a




Figure 1: An example with three paths from the subway to the airport: fastest (solid), suggested by guides (dashed), best for user (dotted).

tourist may only be interested in certain types of venues, such as museums, galleries, and cafes, while certain vehicles may only be able to serve particular customers.

In general, orienteering is an NP-hard problem, and adding categories does not change this fact [2]. We propose an approach based on a best-first strategy to explore the search space, meaning that we first expand the partial solutions that show the greatest potential. We do so with the help of a function approximating the attractiveness of POIs that can still be added. Similar to admissible heuristics in an A*-search, this function needs to satisfy certain properties. Additionally, we are able to prune partial solutions that cannot possibly result in an optimal route.

Even though this technique will speed up the search for an optimal solution by pruning unpromising partial solutions, in the worst case it still has an exponential run time. Therefore, we describe how to transform our method into much more efficient approximation algorithms, creating different variants with important properties concerning the quality of the generated solution and the run time. In summary, we make the following contributions:

- We show how to apply a best-first strategy to the problem of orienteering with category constraints.
- Additionally, we turn the optimal algorithm into different approximation algorithms proving lower bounds

for the quality of a solution and upper bounds for the run time.

- In an experimental evaluation we compare our approach to state-of-the-art algorithms, demonstrating its effectiveness. For some scenarios we improved the quality of the solutions by about 10% with a run time of just five seconds (with the main competitor taking up to 50 seconds and producing worse solutions).

The remainder of the paper is organized as follows. In the next section we cover related work and in Section 3 we formalize the problem. Sections 4 and 5 contain a detailed description of our approach and an approximation algorithm, respectively, while in Section 6 theoretical bounds are provided. This is followed by an experimental evaluation, comparing our algorithm to the state-of-the-art algorithm, a simple greedy heuristic, and an optimal one. Finally, Section 8 concludes the paper.

2. RELATED WORK

A common technique for modeling tourist trip planning [7] and logistics [6] problems is to map them to the *orienteering problem* (OP). Introduced by Tsiligridis in [19], OP is about determining a path from a starting node to an ending node in an edge-weighted graph with a score for each node, maximizing the total score while staying within a certain time budget. Orienteering is an NP-hard problem and algorithms computing exact solutions using branch and bound [8, 13] as well as dynamic programming techniques [11, 15] are of limited use, as they can only solve small problem instances. Consequently, there is a body of work on approximation algorithms and heuristics, most of them employing a two-step approach of partial path construction [9, 19] and (partial) path improvement [1, 4, 16]. Meta-heuristics, such as genetic algorithms [18], neural networks [20], and ant colony optimization [10] have also been tested. However, none of the approaches investigate OP generalized with categories. For a recent overview on orienteering algorithms, which still mentions orienteering with category constraints as an open problem, see [7].

There is also work on planning and optimizing errands, e.g., someone wants to drop by an ATM, a gas station, and a pharmacy on the way home. The generalized traveling salesman version minimizes the time spent on this trip [14], while the generalized orienteering version maximizes the number of visited points of interest (POIs) given a fixed time budget. However, as there are no scores, no trade-offs between scores and distances are considered.

Another, different take on tourist trip planning is to look at the user experience. For instance, Maruyama et al. investigate feedback elicitation from users in a tourist trip planning context [12], while Castillo et al. look at context-based recommendation and collaborative filtering [3]. However, most of these approaches do not consider the actual route planning or use very simple greedy heuristics.

Adapting an existing algorithm for OP would be a natural starting point for developing an approximation algorithm considering categories. However, many of the existing algorithms have a high-order polynomial complexity or no implementation exists, due to their very complicated structure. Two of the most promising approaches we found were the segment-partition-based technique by

Blum et al. [1] and the method by Chekuri and Pál, exploiting properties of submodular functions [5]. The latter approach, a quasi-polynomial algorithm, is still too slow for practical purposes. Additionally, common to all of the approaches, though, is breaking down the itinerary recursively into smaller and smaller segments, which get assembled into a complete tour. If we just run these algorithms without any alterations on POIs with categories, it is very likely that the solution violates the max-n type constraints. A fix would be to try out all possible distributions of max-n type constraints for every recursive call. For example, given a max-n type constraint of 3 for a category and assuming **recl** computes the left half of a route and **recr** the right one, we would have to make the following calls with max-n type constraints: **recl**(0), **recr**(3); **recl**(1), **recr**(2); **recl**(2), **recr**(1); **recl**(3), **recr**(0). While this would guarantee an answer respecting the max-n type constraints, it would also blow up the computational costs.

Nevertheless, Singh et al. [17] modified the algorithm of Chekuri and Pál by introducing spatial decomposition for Euclidean spaces in the form of a grid, making it more efficient. Our work on CLuster Itinerary Planning (CLIP) [2] for OPs on road networks with category constraints was inspired by [17].

3. PROBLEM FORMALIZATION

We assume a set of points of interest (POIs) $p_i, 1 \leq i \leq n$, represented by \mathbf{P} . The POIs, together with a starting and a destination node, denoted by s and d , respectively, are connected by a complete, metric, weighted, undirected graph $G = (\mathbf{P} \cup \{s, d\}, \mathbf{E})$, whose edges, $e_l \in \mathbf{E} = \{(x, y) \mid x, y \in \mathbf{P} \cup \{s, d\}\}$ connect them. Each edge e_l has a cost $c(p_i, p_j)$ that signifies the duration of the trip from p_i to p_j , while every node $p_i \in \mathbf{P}$ has a cost $c(p_i)$ that denotes its visiting time. Each POI belongs to a certain category, such as *museums*, *restaurants*, or *galleries*. The set of m categories is denoted by \mathbf{K} and each POI p_i belongs to exactly one category $k_j, 1 \leq j \leq m$. Given a p_i , $\text{cat}(p_i)$ denotes the category p_i belongs to and $\text{score}(p_i)$ its score or reward, with higher values indicating higher interest to the user. Finally, users have a certain maximum time in their budget to complete the itinerary, denoted by t_{\max} .

DEFINITION 1. (Itinerary) *An itinerary \mathcal{I} starts from a starting point s and finishes at a destination point d (s and d can be identical). It includes an ordered sequence of connected nodes $\mathcal{I} = \langle s, p_{i_1}, p_{i_2}, \dots, p_{i_q}, d \rangle$, each of which is visited once. We define the cost of itinerary \mathcal{I} to be the total duration of the path from s to d passing through and visiting the POIs in \mathcal{I} , $\text{cost}(\mathcal{I}) = c(s, p_{i_1}) + c(p_{i_1}) + \sum_{j=2}^q (c(p_{i_{j-1}}, p_{i_j}) + c(p_{i_j})) + c(p_{i_q}, d)$, and its score to be the sum of the scores of the individual POIs visited, $\text{score}(\mathcal{I}) = \sum_{j=1}^q \text{score}(p_{i_j})$.*

EXAMPLE 1. *Figure 2 shows an example with four POIs, p_1, p_2, p_3 , and p_4 , along with their distances, visiting times, scores, and categories. We simplify the graph slightly to keep it readable: all POIs of the same category have the same score we also omit some edges. One example itinerary between s and d is the one that includes only p_1 , i.e., $\mathcal{I}_1 = \langle s, p_1, d \rangle$, while a second one includes p_2 and p_3 , i.e., $\mathcal{I}_2 = \langle s, p_2, p_3, d \rangle$, a third itinerary can be $\mathcal{I}_3 = \langle s, p_2, p_3, p_4, d \rangle$. Their costs and scores are as follows:*




Figure 2: Itinerary including $n = 4$ POIs

- $\mathcal{I}_1 = \langle s, p_1, d \rangle$:
 $cost(\mathcal{I}_1) = 4 + 1 + 6 = 11$,
 $score(\mathcal{I}_1) = 0.9$;
- $\mathcal{I}_2 = \langle s, p_2, p_3, d \rangle$:
 $cost(\mathcal{I}_2) = 2 + 1 + 2 + 1 + 3 = 9$,
 $score(\mathcal{I}_2) = 0.5 + 0.9 = 1.4$;
- $\mathcal{I}_3 = \langle s, p_2, p_3, p_4, d \rangle$:
 $cost(\mathcal{I}_3) = 2 + 1 + 2 + 1 + 2 + 1 + 1 = 10$,
 $score(\mathcal{I}_3) = 0.5 + 0.9 + 0.5 = 1.9$.

In general, given a subset P of \mathbf{P} , our goal is to place them in an itinerary from s to d as defined above. Given the traveling and visiting times as well as the scores, we need to build an itinerary with duration smaller than t_{max} and maximum cumulative score. As already mentioned, we introduce an additional constraint specifying the number of POIs per category that can be included in the final itinerary. More precisely, we introduce a parameter max_{k_j} for each category k_j that is set by the user to the maximum number of POIs in a category that he or she prefers to visit during the trip. We are now ready to define the *Orienteering Problem with Maximum Point Categories (OPMPC)*.

DEFINITION 2. (OPMPC) *Given a starting point s , a destination point d , n points of interest $p_i \in \mathbf{P}$, with scores described by the function $score(p_i)$, visiting times $c(p_i)$, $1 \leq i \leq n$, traveling times $c(x, y)$ for $x, y \in \mathbf{P} \cup \{s, d\}$, categories $k_j \in \mathbf{K}$, $1 \leq j \leq m$, and the following two parameters: (a) the maximum total time t_{max} a user can spend on the itinerary and, (b) the maximum number of POIs max_{k_j} that can be used for the category k_j ($1 \leq j \leq m$), a solution to the OPMPC is an itinerary $\mathcal{I} = \langle s, p_{i_1}, p_{i_2}, \dots, p_{i_q}, d \rangle$, $1 \leq q \leq n$, such that*

- the total score of the points, $score(\mathcal{I})$, is maximized;
- no more than max_{k_j} POIs are used for category k_j ;
- the time constraint is met, i.e., $cost(\mathcal{I}) \leq t_{max}$.

EXAMPLE 2. *In the presence of categories k_1 with $max_{k_1} = 1$ and k_2 with $max_{k_2} = 1$, and assuming that $t_{max} = 10$, we can observe the following about the itineraries in Example 1: Itinerary \mathcal{I}_1 is infeasible since its cost is greater than t_{max} , while the other two fulfill the time requirement. Comparing \mathcal{I}_2 and \mathcal{I}_3 , we can see that \mathcal{I}_3 is of higher benefit to the user, even though it takes more time to travel between s and d . However, it cannot be chosen since it contains two POIs from k_2 . Itinerary \mathcal{I}_2 contains two POIs, each from a different category and it could be one recommended to the user.*

In Table 1 we give a summary of the notation we use in this paper.

p_i	a single POI
\mathbf{P}	the set of all POIs
n	number of POIs
k_j	a single category
\mathbf{K}	the set of all categories
m	number of categories
max_{k_j}	category constraint for category k_j
$cat(p_i)$	category of POI p_i
$score(p_i)$	score of POI p_i
$c(p_i)$	cost of p_i (visiting time)
$utility(p_i)$	utility of p_i (used for greedy strategy)
G	graph containing POIs as nodes (basically a road network)
$c(p_i, p_j)$	cost of an edge
\mathcal{I}	itinerary
t_{max}	time constraint
s	starting node
d	destination node
u_{k_j}	no of POIs of category k_j in the itinerary
$\mathbf{P}_{rem}(\mathcal{I})$	set of POIs \mathbf{P} without POIs in \mathcal{I}
$cost(\mathcal{I})$	total cost of \mathcal{I} (traveling and visiting times)
$score(\mathcal{I})$	current score of an itinerary
$pot(\mathcal{I})$	potential score of an itinerary (upper bound on achievable score)
$extra(\mathcal{I})$	upper bound on score of still reachable POIs in $\mathbf{P}_{rem}(\mathcal{I})$

Table 1: Notation

4. OUR APPROACH

We roughly follow a best-first strategy, meaning that we keep all solutions generated so far in a priority queue sorted by their potential score. A solution is represented by the set of POIs it contains. In each step, we take the solution with the highest potential score from the queue, expand it with POIs that have not been visited yet, and re-insert the expanded solutions back into the queue. An important difference to the classic best-first search is that in our case it is not straightforward to identify the goal state, as we do not know the score of the optimal itinerary a priori. Consequently, we have to run the algorithm until there are no solutions left in the queue that have a higher potential score than the best found so far (which then becomes the result we return). Applying this approach in a straightforward fashion is not practical in most cases as OPMPC is an intractable problem. Therefore, after presenting an algorithm that computes the optimal solution, we turn it into an approximation algorithm.

4.1 Potential Score

An important aspect of our search strategy is the computation of the potential score of a partial solution, which has to be an upper bound of the score of the fully expanded solution based on this partial solution. Basically, this follows the principle of admissible heuristics found in the A*-algorithm, with the difference that we never underestimate the value, since we consider scores and not costs.

DEFINITION 3. (Potential score) *Given a (partial) itinerary $\mathcal{I} = \langle s, p_1, p_2, \dots, p_i, d \rangle$, its potential score $pot(\mathcal{I})$ is defined as follows:*

$$pot(\mathcal{I}) = score(\mathcal{I}) + extra(\mathcal{I})$$

where $\text{score}(\mathcal{I}) = \sum_{j=1}^i \text{score}(p_j)$ and $\text{extra}(\mathcal{I})$ is a heuristic never underestimating the additional score that is still possible for \mathcal{I} .

Algorithm 1 shows the pseudocode for computing $\text{extra}(\mathcal{I})$. We obtain an approximation (without underestimating the true extra score we can still get) by figuring out how many of the top-scoring POIs we could potentially still fit into the partial itinerary \mathcal{I} . In a first step, we remove all POIs that are already in \mathcal{I} from the set of all POIs (line 2). We call this set of remaining POIs $\mathbf{P}_{\text{rem}(\mathcal{I})}$. Then we discard all POIs in $\mathbf{P}_{\text{rem}(\mathcal{I})}$ that are too costly to travel to from the penultimate POI p_i to d (line 4). Now that we have reduced the set of POIs, we turn to the category constraints to get an even more accurate picture (line 6). For every category we know how many POIs we can still fit into the itinerary without violating the category constraints. As previously defined, \max_{k_j} is the maximum number of POIs that can be chosen for category k_j . Let u_{k_j} be the number of POIs of category k_j currently found in \mathcal{I} (line 7), then we know that we can only fit $(\max_{k_j} - u_{k_j})$ more POIs of category k_j into \mathcal{I} (line 8). For each category, we sort the POIs in $\mathbf{P}_{\text{rem}(\mathcal{I})}$ according to their score in descending order and add up the scores of the top $(\max_{k_j} - u_{k_j})$ POIs of each category (lines 10 to 13). This comprises $\text{extra}(\mathcal{I})$.

Algorithm 1: $e \leftarrow \text{Extra}(q, \mathcal{I}, G, \mathbf{P})$

Input: query q (consisting of s, d, t_{\max} , and category constraints), an itinerary \mathcal{I} , graph G , set \mathbf{P} of all POIs

Output: potential extra score $\text{extra}(\mathcal{I})$

```

1  $x \leftarrow 0$ 
2  $\mathbf{P}_{\text{rem}(\mathcal{I})} \leftarrow \mathbf{P} \setminus \mathcal{I}$ 
3  $p_i \leftarrow$  last but one node in  $\mathcal{I}$ 
4 foreach  $p \in \mathbf{P}_{\text{rem}(\mathcal{I})}$  do
5   if  $\text{cost}(\mathcal{I}) - c(p_i, d) + c(p_i, p) + c(p) + c(p, d) > t_{\max}$ 
     then  $\mathbf{P}_{\text{rem}(\mathcal{I})} \leftarrow \mathbf{P}_{\text{rem}(\mathcal{I})} \setminus \{p\}$ 
6 foreach category  $k_j \in q$  do
7    $u_{k_j} \leftarrow |\{p \in \mathcal{I} | \text{cat}(p) = k_j\}|$ 
8   if  $\max_{k_j} > u_{k_j}$  then
9      $\mathbf{P}_{\text{rem}(\mathcal{I})}^{k_j} \leftarrow \{p \in \mathbf{P}_{\text{rem}(\mathcal{I})} | \text{cat}(p) = k_j\}$ 
10    sort  $\mathbf{P}_{\text{rem}(\mathcal{I})}^{k_j}$  in descending order of scores
11    for  $i$  from 1 to  $\max_{k_j} - u_{k_j}$  do
12       $p_i \leftarrow i$ th POI in  $\mathbf{P}_{\text{rem}(\mathcal{I})}^{k_j}$ 
13       $x \leftarrow x + \text{score}(p_i)$ 
14 return  $x$ 

```

4.2 Our Algorithm

We are now ready to describe the pseudocode in Algorithm 2. For the moment disregard the parts marked in gray: these are optimizations that will be discussed in the following section. Partial solutions are kept in a double-ended priority queue (also called deque). The deque is initialized with the empty solution e just containing s and d and a score of 0. We take the solution with the highest potential score out of the deque and expand it, one by one, with the POIs that have not been visited yet. In order to figure out whether a solution violates the time constraint t_{\max} , we

have to compute the (shortest) length of the itinerary of a solution. As every partial solution contains only a limited number of POIs, we compute this length in a brute-force way. We discard any expanded solutions that violate the time constraint or any of the category constraints and put all valid expansions back into the deque. If an expanded solution is better than the current best solution, we update it. We continue until the potential score of the solution taken from the deque is smaller than the best found so far (which then becomes the answer we return).

Algorithm 2: $b \leftarrow \text{MainLoop}(q, G)$

Input: query q (consisting of s, d, t_{\max} , and category constraints), graph G , set \mathbf{P} of all POIs

Output: best solution b

```

1  $Q$  priority deque for partial solutions
2  $e.\text{cost} \leftarrow c(q.s, q.d)$ 
3  $e.\text{extra} \leftarrow \text{Extra}(q, e, G, \mathbf{P})$ 
4  $b \leftarrow \text{ExtendGreedy}(q, e)$ 
5  $Q.\text{push}(e)$ 
6 while not  $Q.\text{empty}()$  do
7    $s \leftarrow Q.\text{popMaximum}()$ 
8   foreach  $z \in \text{Expand}(s, G)$  do
9     if  $z$  violates any category constraint then
       continue
10     $z \leftarrow \text{ComputePathAndScore}(z)$ 
11    if  $z.\text{cost} > q.t_{\max}$  then continue
12     $z.\text{extra} \leftarrow \text{Extra}(q, z, G, \mathbf{P})$ 
13    if  $z.\text{score} > b.\text{score}$  then  $b \leftarrow z$ 
14     $Q.\text{push}(z)$ 
15    if at an early stage then
16       $y \leftarrow \text{ExtendGreedy}(q, z)$ 
17      if  $y.\text{score} > b.\text{score}$  then  $b \leftarrow y$ 
18    while  $b.\text{score} > Q.\text{minimum}().\text{score} + Q.\text{minimum}().\text{extra}$  do
19       $Q.\text{popMinimum}()$ 
20 return  $b$ 

```

4.3 Further Optimizations

With the help of the score of the current best solution, we can do some pruning, because only partial solutions whose potential score is better have to be considered. In order to be able to prune earlier and more aggressively, we initialize the best solution found so far with an itinerary found quickly using a greedy algorithm (see line 4 and Algorithm 3). Given a partial itinerary $\mathcal{I} = \langle s, p_1, p_2, \dots, p_i, d \rangle$, the greedy strategy adds to this path a POI $p \in \mathbf{P}_{\text{rem}(\mathcal{I})}$ such that its utility

$$\text{utility}(p) = \frac{\text{score}(p)}{c(p_i, p) + c(p) + c(p, d)}$$

is maximal and no constraints are violated. We repeat this until no further POIs can be added to the itinerary.

We can also use a greedy approach to improve the current best solution while the algorithm is running (see lines 15 to 17). When expanding a partial solution taken from the deque, we also complete it with a greedy strategy. This improves the current best solution more rapidly, leading to more aggressive pruning and is especially useful in earlier

Algorithm 3: $s \leftarrow \text{ExtendGreedily}(q, z)$

Input: query q , solution z
Output: solution s

```
1  $s, s' \leftarrow z$ 
2 while  $\text{length}(s') \leq q.t_{\max}$  do
3    $s \leftarrow s', u \leftarrow -\infty$ 
4   foreach  $p \in \mathbf{P}_{\text{rem}(s, \mathcal{I})}$  do
5     if  $p$  violates any category constraint of  $s'$  then
6       continue
7     if  $u < \text{utility}(p)$  then
8        $u \leftarrow \text{utility}(p), p' \leftarrow p$ 
9   AppendToEnd( $s', p'$ )
10 return  $s$ 
```

stages of the algorithm, as there is still a lot of unexplored search space. We may want to stop the greedy expansion in later stages, due to it being less effective, as we have already explored a lot of promising directions. The threshold $g \in [0, 1]$ defines when the algorithm is in a *early stage*, it checks each expanded solution z : whenever the ratio $z.\text{cost}/t_{\max} \leq g$, it does a greedy expansion. We investigate this parameter in more detail in Section 7 on the experimental evaluation.

5. APPROXIMATION ALGORITHMS

Even though pruning and further optimizations speed up the search for a solution, in the worst case we still face exponential run time. After all, we are looking for an optimal solution for an NP-hard problem. Here we present variations of Algorithm 2 that turn it into an approximation algorithm. In principle, we can guarantee an upper bound on the run time or a lower bound on the score.

5.1 Bounding the Score

In line 18 of Algorithm 2 we remove (partial) solutions from the end of the deque that have a potential score smaller than the best solution b found so far. Since solutions are sorted by an overestimation of their score, we never accidentally get rid of a partial solution that could be expanded into the optimal one. However, we could prune more aggressively by introducing a *cut factor* c ($c \geq 1$) and checking whether $c \cdot b.\text{score} > Q.\text{minimum}().\text{score} + Q.\text{minimum}().\text{extra}$.¹ The larger c , the more partial solutions will be ignored. However, by doing so we also risk losing the optimal and other good solutions. Nevertheless, we are sure to get a solution that guarantees at least $\alpha = 1/c$ of the score of the optimal solution.

5.2 Bounding the Run Time

While the approach in the previous section makes sure that we always get a certain amount of the optimal score, we are not able to bound its run time. All we know is that it is faster, since we prune more partial solutions. By modifying Algorithm 2 in a different way, we can obtain a run time bound. Instead of allowing the queue to become arbitrarily long, we limit its length to a maximum of l_{\max} entries. In this way, we put a limit on the number of partial solutions that can be expanded. Before pushing an expanded solution

¹If we set c to 1, we get the original algorithm, i.e., the approximation algorithm is actually a generalization.

z back into the deque in line 14, we check whether there is still space. If there is not enough space, we remove the partial solution with the lowest potential score (this could also be z). While we cannot determine the quality of the answer a priori, we know the value for α upon completion of the algorithm if we keep track of the largest potential score of all the partial solutions we discarded. Assume that z_{\max} has the largest potential score among the discarded solutions. Then we know that $\alpha = b/\text{pot}(z_{\max})$, where b is the optimal solution.

Another, more direct, way of controlling the run time is to set an explicit limit r for it. After the algorithm uses up the allocated time, it returns the best solution found so far. For example, we can check between line 7 and 8 whether there is any time left. If not, we jump out of the while-loop. Again, it is not possible to prescribe a value for α beforehand, but we can determine its value when the algorithm finishes: in this case z_{\max} is the partial solution with the largest potential score that is still in the deque when we stop.

6. PROPERTIES AND BOUNDS

In the following we prove that the pruning we do leads to an algorithm computing the itinerary with the maximal score. Additionally, we prove the bounds for the score and the run time of the approximation algorithms.

6.1 Correctness of Pruning

Let \mathcal{I}_c be a *complete itinerary*, i.e., no further POIs can be added to \mathcal{I}_c without violating a constraint. We have to prove that for every complete itinerary \mathcal{I}_c , $\text{pot}(\mathcal{I}_c) = \text{score}(\mathcal{I}_c) \leq \text{pot}(\mathcal{I}_p)$, where \mathcal{I}_p is a partial itinerary of \mathcal{I}_c ($\mathcal{I}_p \subset \mathcal{I}_c$). $\mathcal{I}_p \subset \mathcal{I}_c$ iff the POIs p_1, p_2, \dots, p_p in \mathcal{I}_p are a subset of the POIs p_1, p_2, \dots, p_c in \mathcal{I}_c . The potential score of \mathcal{I}_c is equal to its score, since $\text{extra}(\mathcal{I}_c) = 0$ (no more POIs can be added). We can actually show a stronger statement that includes the one above.

LEMMA 1. *As we expand solutions, their potential score decreases: given two itineraries \mathcal{I}_1 and \mathcal{I}_2 with $\mathcal{I}_1 \subset \mathcal{I}_2$, it follows that $\text{pot}(\mathcal{I}_2) \leq \text{pot}(\mathcal{I}_1)$.*

PROOF. We prove the lemma by structural induction.

- **Induction step** ($n \rightarrow n+1$):

Given two itineraries \mathcal{I}_n with POIs p_1, p_2, \dots, p_n and \mathcal{I}_{n+1} with POIs p_1, p_2, \dots, p_{n+1} , we have to show that $\text{pot}(\mathcal{I}_{n+1}) \leq \text{pot}(\mathcal{I}_n)$

$$\begin{aligned} &\Leftrightarrow \text{score}(\mathcal{I}_{n+1}) + \text{extra}(\mathcal{I}_{n+1}) \\ &\leq \text{score}(\mathcal{I}_n) + \text{extra}(\mathcal{I}_n) \\ &\Leftrightarrow \text{score}(\mathcal{I}_n) + \text{score}(p_{n+1}) + \text{extra}(\mathcal{I}_{n+1}) \\ &\leq \text{score}(\mathcal{I}_n) + \text{extra}(\mathcal{I}_n) \\ &\Leftrightarrow \text{score}(p_{n+1}) + \text{extra}(\mathcal{I}_{n+1}) \\ &\leq \text{extra}(\mathcal{I}_n) \end{aligned}$$

Assuming that we can still add (up to) k POIs to \mathcal{I}_n , $\text{extra}(\mathcal{I}_n)$ is computed by taking the top- k POIs from $\mathbf{P}_{\text{rem}(\mathcal{I}_n)}$, whereas $\text{extra}(\mathcal{I}_{n+1})$ is computed by taking the top- $(k-1)$ POIs from $\mathbf{P}_{\text{rem}(\mathcal{I}_{n+1})}$. We have just expanded \mathcal{I}_n by adding p_{n+1} to it, so we know that $\mathbf{P}_{\text{rem}(\mathcal{I}_{n+1})} \cup \{p_{n+1}\} \subseteq \mathbf{P}_{\text{rem}(\mathcal{I}_n)}$, as p_{n+1} is missing from $\mathbf{P}_{\text{rem}(\mathcal{I}_{n+1})}$ (compared to $\mathbf{P}_{\text{rem}(\mathcal{I}_n)}$) and

$\mathbf{P}_{rem(\mathcal{I}_{n+1})} \cup \{p_{n+1}\}$ can at most reach $\mathbf{P}_{rem(\mathcal{I}_n)}$, because the remaining time in itinerary \mathcal{I}_{n+1} is less than the remaining time in itinerary \mathcal{I}_n . Therefore, the choice we get when picking top-scoring POIs we pick for the left-hand side of the inequality is more restricted compared to the the right-hand side.

- **Base case (empty itinerary):**

The empty itinerary \mathcal{I}_0 contains no POIs, therefore $\text{pot}(\mathcal{I}_0) = \text{extra}(\mathcal{I}_0)$ and we choose the top-scoring POIs from the largest possible set $\mathbf{P}_{rem(\mathcal{I}_0)}$, which means that \mathcal{I}_0 has the largest possible potential score.

□

6.2 Lower Bounding the Score

Using the approximation technique from Section 5.1, we can guarantee a lower bound for the score we achieve. At any given time, when discarding a partial solution \mathcal{I}_p , we know that its potential score $\text{pot}(\mathcal{I}_p)$ is at most a factor c greater than the best solution b found so far: $c \cdot b \geq \text{pot}(\mathcal{I}_p)$. If there is no further change for b , the ratio of achieving $\alpha = 1/c$ of the best score holds, as $\text{pot}(\mathcal{I}_p)$ never underestimates the score of a full expansion of \mathcal{I}_p . If b is replaced by a new best score b' , we then know that the previously discarded partial solutions are actually closer to the optimal solution, as b' can only be greater than b . In fact, the cut factor of the previously discarded partial solutions improves to $c' = b/b' \cdot c$. Additionally, the lemma in the previous section states that the potential score of an expanded solution can never go up, which means that we are also on the safe side here: we cannot lose a complete solution that is less than a factor α of the optimal solution by discarding \mathcal{I}_p .

Even when using the approximation techniques bounding the run time shown in Section 5.2, which do not allow us to set a lower bound for the score a priori, we can draw conclusions about the factor α after the algorithm finishes. When running our algorithm with a limited queue length, we keep track of the solution with the largest potential score $\text{pot}(\mathcal{I}_{\max})$ that we have discarded. When execution stops, we can now compare $\text{pot}(\mathcal{I}_{\max})$ to the best solution b that is returned, knowing that α has to be at least $b/\text{pot}(\mathcal{I}_{\max})$. Using an explicit time limit r , the largest (implicitly) discarded $\text{pot}(\mathcal{I}_{\max})$ can be found at the head of the queue and we can compute α as indicated above. This even makes it possible to run our algorithm in an iterative fashion. Once the time limit r has been reached, but a user is not satisfied with the current level of quality, they can continue running the algorithm for another limited time period, checking α again. In this way, the quality of the solution and the run time can be balanced on the fly.

6.3 Upper Bounding the Run Time

When setting an explicit time limit r , bounding the run time is straightforward. In case of a limited queue length, drawing conclusions about the run time is more complicated. The worst case for a best-first search algorithm is a scenario in which all the POIs have the same score, the same cost, and a similar distance from s , d and each other. In this case we first expand the empty partial solution generating all itineraries of length one. In the next step, we expand all these itineraries to those of length two, then all of those to those of length three, and so on, meaning that no pruning is taking place. We now give an upper bound for the number

of generated itineraries and then illustrate the effect of a bounded queue length on this number.

We have m categories k_j , $1 \leq j \leq m$, each with the constraint \max_{k_j} and the set $K_{k_j} = \{p_i | p_i \in P, \text{cat}(p_i) = k_j\}$ containing the POIs belonging to this category. The set of relevant POIs R is equal to $\bigcup_{j=1, \max_{k_j} > 0}^m K_{k_j}$, we denote its cardinality by $|R|$ (assuming that R only contains POIs actually reachable from s and d in time t_{\max}). For the first POI we have $|R|$ choices, for the next one $|R| - 1$, and so on until we have created paths of length $\lambda = \sum_{j=1, \max_{k_j} > 0}^m \max_{k_j}$. We cannot possibly have itineraries containing more than λ POIs, as this would mean violating at least one of the category constraints. So, in a first step this gives us $\prod_{i=0}^{\lambda-1} (|R| - i)$ different itineraries. This is still an overestimation, though, as it does not consider the t_{\max} constraint, nor does it consider that once an itinerary has reached \max_{k_j} POIs for a category k_j , the set K_{k_j} can be discarded completely when extending the itinerary further. Introducing a maximum queue length of l_{\max} , we arrive at the final number $\prod_{i=0}^{\lambda-1} \min(|R| - i, l_{\max})$.

7. EXPERIMENTAL EVALUATION

We conducted our experiments on a Linux machine (Arch Linux, kernel version 4.8.4) with an i7-4800MQ CPU running at 2.7 GHz and 24 GB of RAM (of which 15 GB were actually used). The algorithms were written in C++, compiled with gcc 6.2 using `-O2`. The priority deque was implemented with an interval heap, the graph representing the city map with an adjacency list. For comparison with the CLIP algorithm we used black-box testing, configuring it with the relaxed linear programming solver as suggested in [2]. The greedy strategy uses Algorithm 3 on an empty itinerary containing just s and d .

We averaged 25 different test runs for every data point in the following plots. For each test run we selected a start and end node randomly, such that $c(s, d) < t_{\max}$. When comparing different algorithms, we generated one set of queries and re-used it for every algorithm. Unless specified otherwise, as default values for the parameters we two hours for t_{\max} , the largest number of categories possible for a graph (four or nine), a category constraint of two for each category, a cut ratio of 1.2, a greedy expansion threshold of 1, an unbounded queue length, and unlimited run time.

7.1 Data Sets

We used several different data sets in our evaluation. The first one was an artificial one consisting of a grid with 10,000 nodes (100×100) with a total of 3,000 POIs in four different categories. The distance between neighboring nodes in the grid was 60 seconds. The POIs were randomly scattered in the grid, had a visiting time between three minutes and one hour and a score between 1 and 100.

The second artificial network is a spider network having the same number of nodes and POIs as the grid, but the edges are placed as a 100-sided polygon with 100 levels. The edges between levels are 100 seconds, and the central, and smallest, polygon has sides of 8 seconds. The other levels become larger and larger as in a Euclidean plane. For a visualization of the grid and spider networks, see Figure 3.

The two real-world data sets were a map of the Italian city of Bolzano with a total of 1830 POIs in nine categories

and a map of San Francisco with a total of 1983 POIs in four categories.



Figure 3: Grid and spider network example

7.2 Effects of Parameters

First of all, we illustrate the effects of the different parameters on the performance of our approximation algorithm. Namely, these are the values for the greedy expansion threshold, the cut ratio, the queue length limit, and the run time limit. We also show how the number of categories selected by a user influences our performance.

7.2.1 Greedy expansion threshold

Figure 4 shows the impact of the greedy expansion threshold on the performance of our algorithm (here on a grid network; the results for the other networks look very similar). As a reminder, the threshold g determines when we stop expanding a partial solution greedily. If the accumulated cost of an itinerary so far divided by t_{\max} is less than or equal to g , we do the expansion ($g = 0$ meaning we never expand, $g = 1$ meaning we always do so).

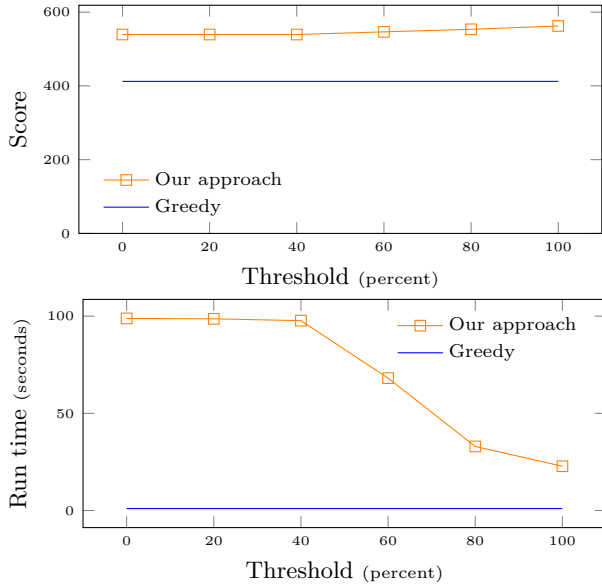


Figure 4: Varying the greedy expansion threshold

As can be seen in Figure 4 (top), the effect on the score is minimal (for comparison a pure greedy strategy is also shown). However, the improvements in terms of run time are significant (Figure 4, bottom). Running a greedy expansion during the execution of the algorithm updates the best solution found so far much faster, resulting in better pruning. On average, it was always worth running a greedy expansion right up to the point when the algorithm finishes, since greedily expanding a partial solution can be done very

quickly. In fact, toward the end of the execution this can be done even faster, as only one or two more POIs can be added to an itinerary.

7.2.2 Cut ratio

Next, we investigate the impact of the cut ratio. Increasing the cut ratio allows us to prune more aggressively, albeit at the price of sometimes losing the optimal solution. In Figure 5 we depict the results for varying the cut ratio from 1 (optimal case) to 2.5 (guaranteeing at least 0.4 of the optimal score). As expected, the score of the found solutions goes down when increasing the cut ratio. This is the case for all data sets. Nevertheless, it does so much slower than the guaranteed lower bound for the score and it also stays well above the score for the solutions found by the simple greedy algorithm. This is especially true for the real-world data sets (Figures 5(c), (d), (e), and (f)). We found that our algorithm finds good solutions early (more on this in Section 7.2.4 about setting a fixed run time limit).


The impact of the cut ratio on the run time is much larger, though (see Figure 6). For some data sets, namely the grid and San Francisco networks, increasing the cut ratio brings down the run time significantly.² Grid networks, including the artificial grid and San Francisco, whose map is composed of many grid-like structures, have a very long run time for small cut ratios, while spider networks do not exhibit this behavior. (Bolzano, which has grown organically over centuries around an old city center, shares some of the features of a spider network.) The different edge lengths make the central part of the spider network much easier to explore, whereas the external parts are more isolated. As a consequence, the computation of $\text{extra}(\mathcal{I})$ is more precise, meaning that we get a much more accurate picture about the potential scores of the yet unexpanded solutions that are still in the queue, making it easier to prune partial solutions. In a grid-like network, there are lots of alternatives, making it more difficult to compute $\text{extra}(\mathcal{I})$ accurately, so the potential scores of some unexpanded solutions may still be high, and we need more time to check them.

7.2.3 Queue length


Now we look at the impact of bounding the queue length to shorten the run time of our algorithm. The left-hand column of Figure 7, labeled (a), shows the results for the grid network, the right-hand column, labeled (b), those for the spider network. The cut ratio is set to 1 to measure only the impact of the queue length limit.³ The size of the queue on the x-axis is measured relatively to the number of POIs in the graph and the horizontal line at the top of the figure is the optimal result when running our algorithm with an unbounded queue. As can be seen in the top row of Figure 7, having a very short queue lowers the score slightly. We do not start the plot y-axis at 0 to make the small differences of the score more visible. That is also the reason why we did not add the score of the greedy algorithm (it is 412 for Figure 7(a) and 252 for Figure 7(b)). Unsurprisingly, the run time goes up with increasing queue length (see bottom

²The two hour itineraries in the real-world setting for Bolzano and San Francisco are a bit too short to notice an effect.


³Due to the cut ratio, we did not run this experiment for the real-world data sets, as it would have taken too much time.




(a) Grid network




(b) Spider network




(c) Bolzano network ($t_{\max} = 2h$)



(d) San Francisco network ($t_{\max} = 2h$)



(e) Bolzano network ($t_{\max} = 3h$)



(f) San Francisco network ($t_{\max} = 3h$)

Figure 5: Increasing the cut ratio




Figure 6: Effect of the cut ratio on run time

row of Figure 7) and it will eventually reach the run time of the unbounded queue.

7.2.4 Run Time Limit

Bounding the queue length is an indirect way of controlling the run time, in Figure 8 we show results for setting an explicit run time limit. The orange line at the top of of each plot represents the optimal score. We can see that even for very short run time limits we get results for every data set that are very close to the optimal score. What this

means is that our algorithm finds very good solutions early in the search process and then, if we let it continue, spends the remaining time basically verifying that there are not significantly better solutions. We also clearly outperform the greedy heuristic in terms of score (the greedy algorithm takes between 0.2 and 1.5 seconds to find a solution).

When limiting the run time, we cannot guarantee the quality of the attained score a priori. Nevertheless, when the algorithm stops running, together with the score it returns the factor α , that tell us which proportion of the optimal

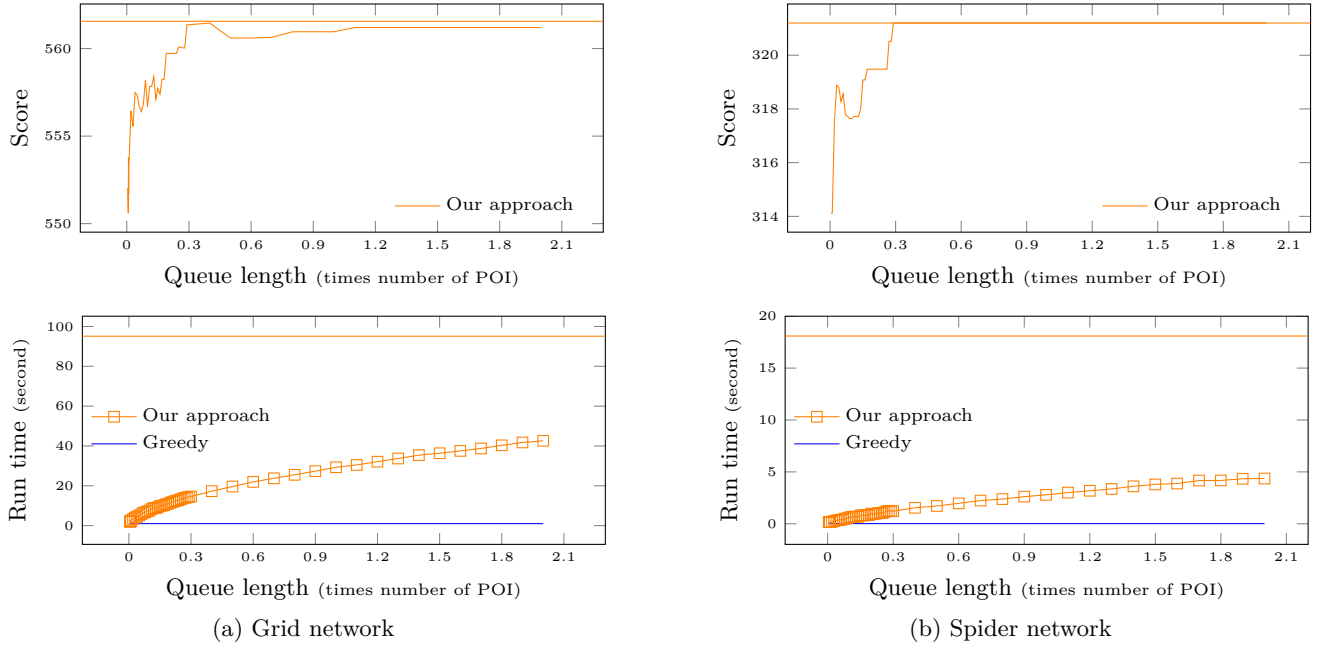


Figure 7: Limiting the queue length

score we have reached at least (for details, see Section 6.2). We have plotted the lower bound α in Figure 8, please note that for this measure the scale on the right-hand side of the y-axis is used. This has important implications for running our algorithm: we basically have an anytime algorithm. We can stop it at any time and get a solution and its quality. Eventually we continue with the execution and investigate the current state a short time later, repeating the process. Consequently, we are able to balance the quality of the solution and the run time on the fly.

Looking at Figure 8, we see that the results for two of the data sets, the spider network in Figure 8(b) and the Bolzano network in Figure 8(c) are somewhat different, which comes as a surprise, as we mentioned that the Bolzano network shares some of the properties of a spider network in Section 7.2.2. However, there are other effects at work here, too. For the Bolzano network the larger number of categories is responsible for the lower guarantee. This makes it more difficult to compute $\text{extra}(I)$ accurately, making it more difficult to discard potential unexpanded solutions. (We go into more details about the number of categories in the following section.) We would like to point out that the found solutions are close to the optimal for every data set, only the lower bounds for the score are affected differently.

7.2.5 Impact of Categories

For the final parameter, we illustrate the impact of categories on the run time. While our approximation algorithms scale well in terms of the itinerary length (t_{\max}), as we will see in Section 7.3, this is not independent of the number of categories. Figure 9(a) shows the effects of increasing the number of categories. By putting a limit on the run time, we can keep the execution time of the algorithm low, while still generating very good solutions. In the upper diagram, we only show the plots for two and eight categories to keep it uncluttered. In the lower diagram, the run time for two categories actually goes down for longer itineraries. At the

end of visiting two POIs of two categories each, we actually start having more and more spare time, which makes the problem easier to solve.

Figure 9(b) shows the effect of increasing the maximum constraint of categories. Again, we quickly generate good solutions in a scalable way, due to the explicit limit on the run time.

7.3 Comparison with Competitors

Here we compare our algorithm to the state-of-the-art. We distinguish two different cases: (a) the artificial networks (grid and spider) and (b) the real-world data sets (Bolzano and San Francisco). For the real-world data sets we also run more realistic queries with a larger time constraint t_{\max} and for that reason do not compare it to the optimal algorithm (as its run time explodes for long itineraries).

7.3.1 Grid and Spider Networks

Figure 10 illustrates the results of running our algorithm with a cut factor of 1.2 against the optimal algorithm (the left-hand column, labeled (a), shows the results for the grid network, the right-hand column, labeled (b), shows those for the spider network). In terms of the score our approach outperforms the greedy approach and comes close to the optimal solution, while at the same time having a much better run-time performance than the optimal algorithm. The final measurements for large values of t_{\max} are missing, since we aborted runs taking longer than ten minutes to complete.

7.3.2 Real-world Data Sets

We now move to the real-world data sets. Figure 11(a), found in the left-hand column, compares three different variants of our algorithm (cut factor, bounded queue, and run time limit) with the state-of-the-art algorithm for orienteering with categories, CLIP on a map of Bolzano. In terms of the score (upper part of Figure 11(a)) the three variants

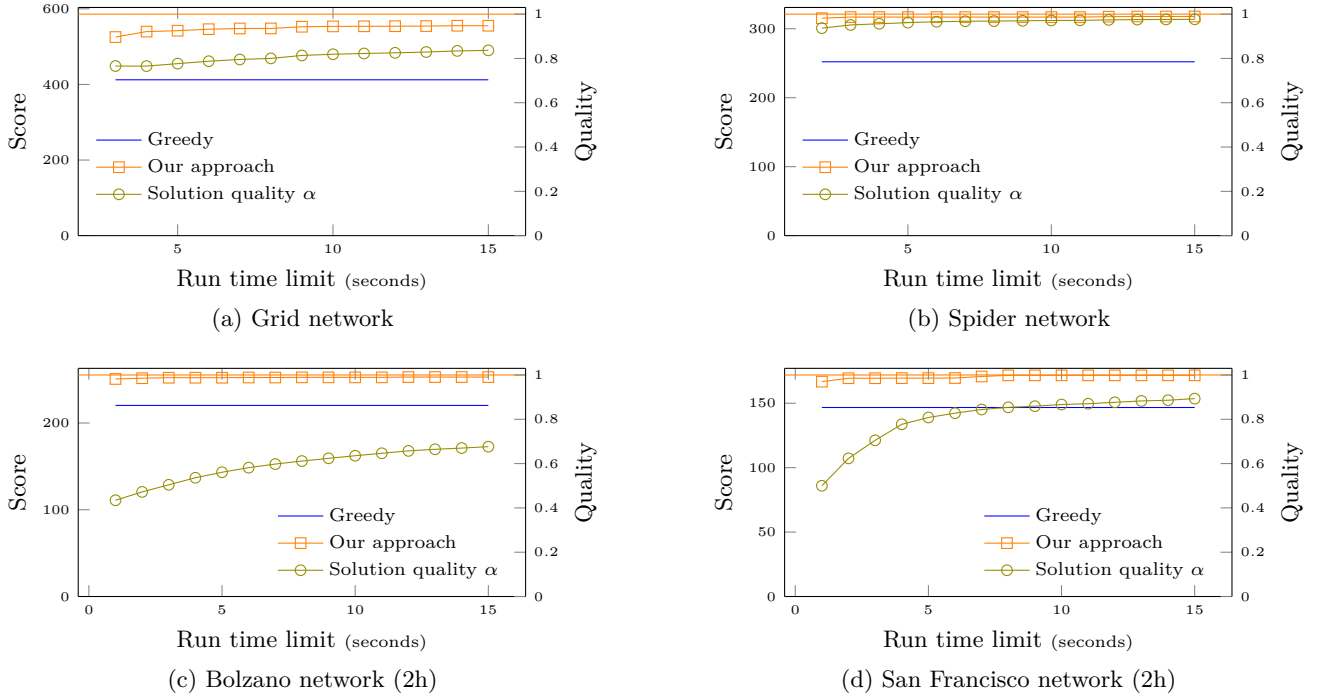


Figure 8: Limiting the run time

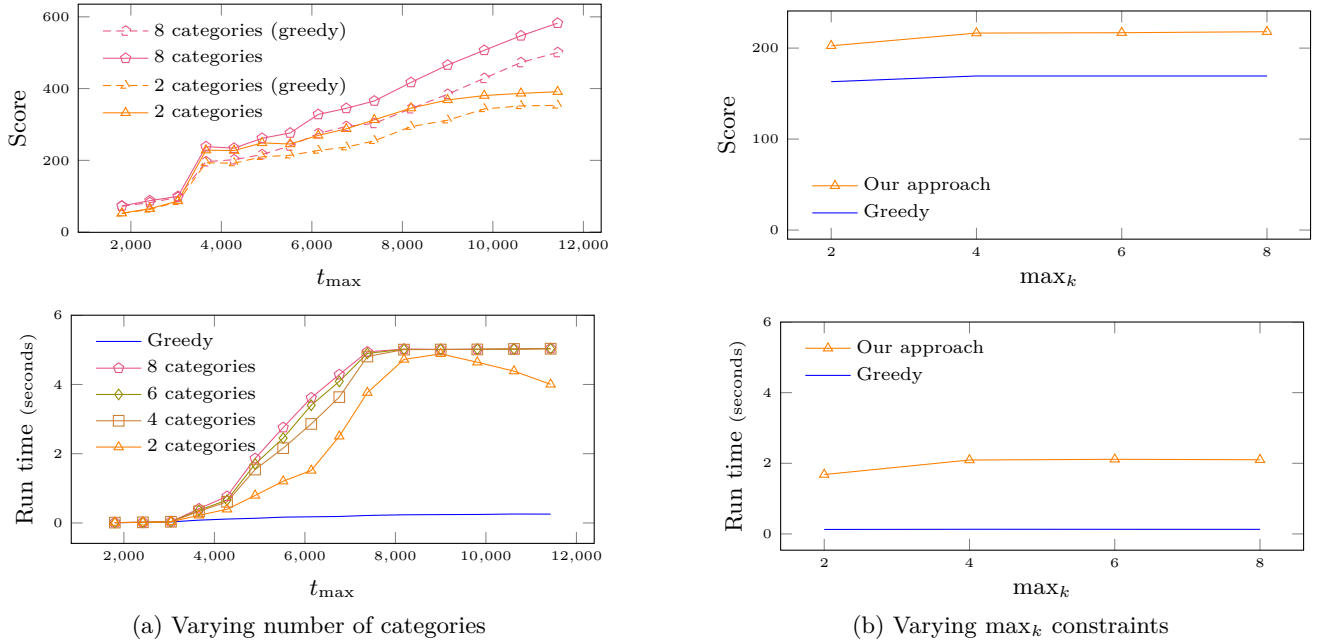




Figure 9: Impact of Categories

of our algorithm are almost identical: merely one measurement resulted in a difference two digits after the decimal point (thus we only depict one curve). Our algorithm outperforms both, CLIP and the greedy heuristic, the latter by a large margin. When we look at the run time (lower part of Figure 11(a)), we see huge differences, though. The only competitive algorithms are the ones limiting the run time in some form, either directly or by limiting the queue length.



In the final experiment we show the full strength of our approach for running queries with a large time constraint

t_{\max} . Here we run a blended version of our algorithm, combining a run time limit of five seconds with a queue length of half the number of total POIs in the graph. At first glance, combining a run time limit with a fixed queue length seems redundant. However, from a certain run time limit onwards, the queue length may become quite large and we may want to restrict it.

Figure 11(b) shows results for the San Francisco data set, comparing CLIP to our blended variant. In terms of the score, our approach outperforms both, CLIP and the greedy





(a) Grid network





(b) Spider network

Figure 10: Comparison with the exact algorithm



(a) Bolzano network



(b) San Francisco network

Figure 11: Comparison with CLIP

heuristic. However, we get a better score than CLIP with a shorter run time, demonstrating that our technique scales much better than CLIP. In fact, as we were running the experiments on weaker hardware compared to the findings in [2], we could not replicate the ten-hour itineraries for CLIP shown in [2] in reasonable time. This clearly shows that our algorithm is much more suitable for deployment on mobile devices, which have limited computational resources. We also illustrate that limiting the queue length (to a certain extent) does not have adverse effects on the run-time-limited variant.

8. CONCLUSION AND FUTURE WORK

We have developed an effective and efficient approximation algorithm for solving the orienteering problem with category constraints (OPMPC) by applying a best-first strategy, blending it with greedy search, and then limiting its run time. One major advantage of our technique over CLIP, the state-of-the-art approach for OPMPC, is the fact that our technique is an anytime algorithm, which immediately starts to generate solutions with quality guarantees, as it keeps track of potential scores. Consequently, we can run our algorithm for a fixed time or until a certain quality level has been reached, whichever comes first.

Nevertheless, we still see some room for improvement. Profiling our algorithm, we noticed that we spend a considerable amount of time (about 40%) calculating the potential score. If we were able to do this more efficiently, maybe by parallelizing the task, we would be able to create an even more efficient algorithm. For longer itineraries it may also be interesting to move to more efficient techniques for computing itineraries from sets of POIs. On a more general level, our approach could also be viable for other orienteering variants, such as the team orienteering problem or orienteering with time windows.

Acknowledgments

The icons in Figure 1 are from <http://icons8.com> and used under the Creative Commons Attribution-NoDerivs 3.0 Unported (CC BY-ND 3.0) License. To view a copy of the license, visit <https://creativecommons.org/licenses/by-nd/3.0/>.

9. REFERENCES

- [1] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM J. Comput.*, 37(2):653–670, May 2007.
- [2] P. Bolzoni, S. Helmer, K. Wellenzohn, J. Gamper, and P. Andritsos. Efficient itinerary planning with category constraints. In *SIGSPATIAL/GIS'14*, pages 203–212, Dallas, Texas, 2014.
- [3] L. Castillo, E. Armengol, E. Onaindía, L. Sebastián, J. González-Boticario, A. Rodríguez, S. Fernández, J. D. Arias, and D. Borrajo. SAMAP: A user-oriented adaptive system for planning tourist visits. *Expert Systems with Applications*, 34(2):1318–1332, 2008.
- [4] C. Chekuri, N. Korula, and M. Pál. Improved algorithms for orienteering and related problems. In *SODA'08*, pages 661–670, 2008.
- [5] C. Chekuri and M. Pál. A recursive greedy algorithm for walks in directed graphs. In *FOCS'05*, pages 245–253, 2005.
- [6] J. S. Dunn. Scheduling underway replenishment as a generalized orienteering problem. Master's thesis, Naval Postgraduate School, Monterey, California, 1992.
- [7] D. Gavalas, C. Konstantopoulos, K. Mastakas, and G. Pantziou. A survey on algorithmic approaches for solving tourist trip design problems. *Journal of Heuristics*, 20(3):291–328, 2014.
- [8] M. Gendreau, G. Laporte, and F. Semet. A branch-and-cut algorithm for the undirected selective traveling salesman problem. *Networks*, 32(4):263–273, 1998.
- [9] C. Keller. Algorithms to solve the orienteering problem: A comparison. *European Journal of OR*, 41:224–231, 1989.
- [10] Y.-C. Liang, S. Kulturel-Konak, and A. Smith. Meta heuristics for the orienteering problem. In *CEC '02*, pages 384–389, 2002.
- [11] E. H.-C. Lu, C.-Y. Lin, and V. S. Tseng. Trip-mine: An efficient trip planning approach with travel time constraints. In *MDM'11*, pages 152–161, 2011.
- [12] A. Maruyama, N. Shibata, Y. Murata, K. Yasumoto, and M. Ito. A personal tourism navigation system to support traveling multiple destinations with time restrictions. In *18th Int. Conf on Adv. Information Networking and Applications (AINA'04)*, pages 18–21, 2004.
- [13] R. Ramesh, Y.-S. Yoon, and M. H. Karwan. An optimal algorithm for the orienteering tour problem. *INFORMS Journal on Computing*, 4(2):155–165, 1992.
- [14] M. N. Rice and V. J. Tsotras. Parameterized algorithms for generalized traveling salesman problems in road networks. In *SIGSPATIAL/GIS'13*, pages 114–123, Orlando, Florida, 2013.
- [15] G. Righini and M. Salani. Incremental state space relaxation strategies and initialization heuristics for solving the orienteering problem with time windows with dynamic programming. *Computers & OR*, 36(4):1191–1203, 2009.
- [16] Z. Sevkli and F. E. Sevilgen. Variable neighborhood search for the orienteering problem. In *ISCIS'06*, pages 134–143, 2006.
- [17] A. Singh, A. Krause, C. Guestrin, W. J. Kaiser, and M. A. Batalin. Efficient planning of informative paths for multiple robots. In *IJCAI'07*, pages 2204–2211, 2007.
- [18] F. Tasgetiren and A. Smith. A genetic algorithm for the orienteering problem. In *IEEE Congress on Evolutionary Computation*, 2000.
- [19] T. A. Tsiligrides. Heuristic methods applied to orienteering. *J. Operation Research Society*, 35(9):797–809, 1984.
- [20] Q. Wang, X. Sun, B. L. Golden, and J. Jia. Using artificial neural networks to solve the orienteering problem. *Annals of OR*, 61:111–120, 1995.